



March 22, 2010

FEDERAL COMMUNICATIONS COMMISSION
Office of the Secretary
Via ECFS

Re: Comments Regarding the Notice of Proposed Rulemaking in the Matter of Preserving the Open Internet (GN Docket No. 09-191) and Broadband Industry Practices (WC Docket No. 07-52)

Notice of *Ex-Parte* Presentation

On March 22, 2010, on behalf of only myself, I emailed an academic paper to FCC staff Ruth Milkman. The paper may have some ideas useful to the current proceeding as related to wireless networks.

The paper is attached on the following pages.

Sincerely,

A handwritten signature in cursive script that reads "Scott Jordan".

Scott Jordan
Professor of Computer Science
Department of Computer Science
University of California, Irvine
3029 Bren Hall
Irvine, CA 92697
Email: sjordan@uci.edu
Phone: 949-824-2177
Webpage: <http://www.ics.uci.edu/~sjordan>

c.c. Ruth Milkman

Do wireless networks merit different net neutrality than wired networks?

SCOTT JORDAN

University of California, Irvine

1. INTRODUCTION

Net neutrality has typically been discussed in the context of wired networks, with a focus on the wired public Internet. Those who support wired net neutrality generally believe that there is a danger that Internet Service Providers (ISPs) who offer applications may discriminate in favor of themselves over competing application providers. Examples of such conflicts may include cable ISPs that discriminate in favor of their own Voice-over-IP (VoIP) packets over competing VoIP providers' packets and telephone ISPs that discriminate in favor of their own video over IP packets over competing video providers' packets. Those who oppose wired net neutrality generally believe that any such danger does not represent a market failure and that net neutrality regulation will be counterproductive. Good overviews of the arguments on both sides can be found in Clark [2007], Peha [2007], and Jordan [2007].

Recently, the question has arisen over whether and how net neutrality should apply to wireless networks. In wireless networks, similar anticompetitive concerns may apply. Cellular network ISPs may have the incentive to discriminate in favor of their own video packets over competing video providers' packets. In addition, they may have the incentive to discriminate against any applications that compete with their primary revenue streams, including all competing voice and text-messaging applications that run over the Internet Protocol (IP).

In the political arena, those who support wired net neutrality usually support some form of wireless net neutrality (see e.g. Free Press [2010], Google Inc. [2010], New America Foundation et al. [2010], and Center for Democracy & Technology [2010]), and those who oppose wired net neutrality similarly oppose wireless net neutrality (see e.g. AT&T Inc. [2010] and Verizon and Verizon Wireless [2010]). However both sides recognize that compromise may be forthcoming, and both sides acknowledge that wireless networks face greater technical challenges than wired networks, and that wireless ISPs thus need the ability to exercise stronger forms of network management. It remains unclear whether these differences between wired and wireless network technology merit different treatment with respect to net neutrality.

The key question that this paper attempts to address is whether wireless net neutrality should be different than wired net neutrality because of the different technologies used. There is little academic literature that directly addresses net

Author's address: S. Jordan, Dept. of Computer Science, University of California, Irvine, CA 92697. This is a draft, and has not yet been published. This work has been supported by National Science Foundation.

neutrality in wireless networks. Wu [2007] started much of the current debate, where he focussed on whether subscribers should be able to attach wireless devices of their choice. Wu argued for the extension of the Federal Communication Commission’s (FCC) Carterfone rules [FCC 1968] to wireless networks, including a prohibition on locking of devices to a carrier and allowing attachment of compatible and non-harmful devices. To allow such attachment, he proposes that industry or the FCC should define a basic air interface for wireless devices. Wu also argues for the application of net neutrality to wireless networks, which at the time meant application of the FCC’s Internet Policy Statement [FCC 2005], and stated that carriers should meter and charge for bandwidth usage rather than block particular applications. Wu also argued for wireless carrier disclosure of limits, including locks, protocol or application disabling, and bandwidth limits. Finally, he recommended that carriers and equipment manufacturers should work towards standardization of application development platforms.

In response, Hahn et al. [2007] claim that attachment of devices and Quality-of-Service (QoS) are separate issues. Having previously opposed net neutrality as a method to regulate QoS [Hahn and Litan 2007], in this paper they argue against many of Wu’s proposals. First they argue that there is sufficient wireless competition to avoid market failure and that innovation in wireless devices and applications is thriving. Next they provide an economic analysis and argue that the results show that the benefits of device subsidies, device exclusivity, and limits on devices and on applications outweigh the costs of each.

Both of these papers focus primarily on the device attachment issue. Neither focuses on the differences in traffic management between wired and wireless networks, and hence on potential differences with respect to QoS.

The debate has entered the public policy arena. In 2007, Skype Communications submitted a petition [Skype Communications 2007] to the FCC asking them to declare that wireless carrier services are subject to Carterfone rules, including the right to attach non-harmful devices and the right to run applications. They also asked for the FCC to initiate a rulemaking proceeding to determine whether current wireless carrier practices violate these rules, and to create a mechanism to establish technical standards that ensure that non-harmful applications are allowed.

More recently, the FCC, in its *Notice of Proposed Rulemaking on Open Internet Practices* [FCC 2009] (hereafter referred to as the NPRM), asks whether net neutrality should be applied differently to wireless networks than to wired ones. It notes that wireless networks face special challenges due to attenuation and interference, and that they determine how users and devices share scarce resources through control over the frequency, time, and power of wireless devices’ transmissions. It asks whether wireless devices and/or wireless networks merit different treatment.

With respect to device attachment, the NPRM asks whether subscribers should be able to connect wireless devices of their choice providing that they do not “harm the network”. It ponders whether wireless ISPs should allow attachment of any device with a compatible air interface, including tethering, and if so how to prevent harm.

With respect to services, the NPRM considers which applications or services should be covered by a net neutrality requirement. It considers whether to exclude

managed services, defined as IP-based offerings such as voice and subscription video provided over the same networks used for broadband Internet access. It proposes that a nondiscrimination principle should only apply to Internet services, and thus exclude voice and short messaging services (SMS). However, it ponders what to do in 4th generation (4G) wireless networks capable of supporting voice, video, and data services on a converged platform architecture.

With respect to traffic management, the NPRM notes that wireless capacity may be more limited than wired capacity, and that demands can vary dynamically and widely among users. It discusses that wireless networks must be designed to deal with wide variations in signal levels across the service area as well as interference from other devices. It ponders whether wireless networks are more sensitive to user behavior. It asks how these differences in technical characteristics affect the reasonableness of various network management practices, e.g. whether it is reasonable for a wireless ISP to block particular capacity hungry applications. It also asks what impact tethering will have on wireless network congestion, and what network management measures are reasonable in this context.

The key question that this paper attempts to address is whether wireless net neutrality should be different than wired net neutrality because of the different technologies used. We recognize that readers of this paper will have a diverse set of opinions on net neutrality. However, we expect that, independent of one's position on net neutrality, readers may wonder whether wireless networks should be treated similarly or differently than wired networks with respect to net neutrality. That is the central issue of this paper.

The first half of the paper addresses this question without taking a position on net neutrality. However, a few basic hypotheses are necessary to focus the analysis. First, the paper focuses on whether some form of net neutrality is required to ensure a level playing field between application providers who also serve as ISPs and application providers who do not serve as ISPs; other rationales for net neutrality are not considered here. Second, it is assumed that the primary method of potential discrimination is the use of QoS mechanisms such as packet prioritization or bandwidth reservation.

The primary focus is thus placed on applications and traffic management, rather than device attachment. We are concerned with which applications or services should be covered by a net neutrality requirement, and whether this requires the definition of *managed services*. We are concerned with whether the challenges of wireless signals and mobility merit different traffic management techniques, and how these techniques may affect net neutrality.

There seem to be several parts of this issue. First, whatever one's position on net neutrality in wired networks, do the differences in the definition of *reasonable network management* or *managed services* between wired and wireless networks merit *different treatment* with respect to net neutrality? The resolution of this issue seems to depend in part on whether wireless networks require qualitatively different *types of network management* than wired networks. Second, should wireless network operators have a different ability to *restrict applications* used on their networks than wired ISPs? The resolution of this issue seems to depend in part on whether applications have a greater ability to negatively interfere with desired net-

work operation in wireless networks than in wired networks. Third, should wireless network operators have a different ability to *restrict devices* used on their networks than wired ISPs? The resolution of this issue seems to depend in part on whether wireless devices have a greater ability to negatively interfere with desired network operation than their wired counterparts.

The focus in this paper is on the first two issues (reasonable network management and restricting applications); the third issue (restricting devices) requires additional consideration of interconnection policy that is outside the scope considered here. We focus here on the technical aspects of these issues. With respect to all three issues, there are clearly economic and legal aspects that should be considered. For instance, competition may be substantially different in wireless networks than in wired networks. However, these economic and legal aspects are beyond the scope of this paper.

Section 2 reviews the pertinent aspects of network architecture and discusses the main technical differences between wired and wireless networks. Section 3 more specifically addresses how wired and wireless networks differ with respect to traffic management. With that basic understanding of network architecture, in section 4 we turn to the question of how the differences in traffic management affect net neutrality. These three sections are written without taking a position on net neutrality.

The remaining sections of the paper consider potential methods to extend wired net neutrality to wireless networks. In these sections, we do take a position. In section 5, we state our general position on net neutrality, and then examine how it may or may not apply to wireless net neutrality. Section 6 focuses on how pricing can be used to define a contract with users that removes the need for network control over user applications. Finally, section 7 proposes wireless network net neutrality statute language.

2. HOW ARE WIRELESS NETWORKS DIFFERENT THAN WIRED NETWORKS?

Telephone networks, cable video networks, wireless networks, and the Internet are all based on the concept of a layered architecture. Each network device, and the network as a whole, is abstractly modeled as being composed of a number of vertical layers. Each layer provides certain functionalities. Layering is a form of modularity. In modular design, a designer of a module need only understand the functionality and the interface, not the detailed operation, of other interoperating modules. Although designing a component in a modular fashion restricts the design space, the benefits usually outweigh the costs.

The reference model for layered architectures is the OSI model, developed by the International Standards Organization. The OSI model is composed of 7 layers, as pictured in figure 1. It is useful to think of the physical connection, e.g., wire, as being located below the bottom-most layer (layer 1) and the user, e.g., you, as being located above the top-most layer (layer 7).

OSI layer 1, called the physical layer, implements encodes a bit into a physical signal and vice versa. OSI layer 2, the data link layer, translates a packet into a set of bits and vice versa, and implements a set of rules (called a protocol) about which device can transmit when. OSI layer 3, the network layer, is concerned with

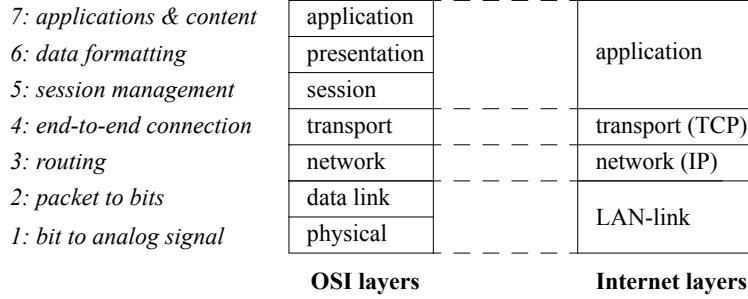


Fig. 1. OSI and Internet layered models

routing a packet from one network device to the next. OSI layer 4, the transport layer, is concerned with functionality required to form a complete connection between a source and destination, including dealing with lost packets and responding to congestion. OSI layer 5, the session layer, manages an entire communication session, e.g., logging onto a service. OSI layers 6, the presentation layer, concerns data presentation, e.g., file or video compression. Finally, OSI layer 7, the application layer, deals with user applications and other high-level functionality, e.g., web browsing, e-mail, file transfer, file sharing, instant messaging, gaming, etc.

Although the OSI model serves as a reference for all network architectures, different networks have modified the model for their own use. As an example, the Internet uses a model with a reduced number of layers, as pictured in figure 1 [Braden 1989]. OSI layers 1 and 2 are combined into a single Internet LAN-link layer. OSI layer 3 is also called the Internet network layer; it includes the Internet Protocol (IP). OSI layer 4 is also called the Internet transport layer; it includes the Transmission Control Protocol (TCP). OSI layers 5 through 7 are merged into a single Internet applications layer.

We will consider differences between wireless and wired networks by the layer, starting at the bottom. The design of the OSI layer 1 protocol of a network is very tightly connected to the type of transmission medium used in the network, e.g. wireless, fiber, coaxial copper, etc.. Each transmission medium has different characteristics of how signals propagate through or down the medium. There are two key physical layer challenges.

First, signals become weaker (or *attenuate*) as they propagate. Wireless signals usually propagate in free space, and hence the signal energy is spread out over all directions in three dimensional space, quickly growing weaker. In contrast, wired signals propagate down a guided path, and hence the signal energy does not spread out as much. Attenuation is thus usually a greater challenge in wireless networks. Second, noise and interference cause difficulties. The environment adds background noise to the signal. In addition, a wireless signal bounces off of many objects (e.g. buildings), and these additional copies of the signal (called *multipath*) can be either beneficial or detrimental. If a wireless user is also mobile, then the amount of attenuation and multipath are continually changing.

The OSI layer 1 protocol of a network is thus matched to the characteristics of the transmission medium used in the network; different mediums warrant different

LAN-link layer protocols. Signals are encoded (or *modulated*) in a manner that is effective given the attenuation and noise challenges. Wireless networks often use more complex modulation schemes than wired networks, in order to combat more severe attenuation and interference. In addition, the modulation scheme determines in what dimension users share network capacity¹. It may specify that a transmission use all of the available frequency or only a specified portion (or *channel*) of the available frequency. Many wireless modulation schemes intentionally encourage multiple transmissions to talk over each other by simultaneously using all available frequency, and the scheme sorts it out later. Finally, wireless modulation schemes continually monitor and modify the transmission power levels of each user (i.e. telling each network device to talk more softly or more loudly) to deal with changing attenuation and multipath levels.

We now move up to OSI layer 2, the data link layer. The design of the OSI layer 2 protocol of a network is also tightly connected to the type of network. The data link layer protocol, given the modulation scheme implemented in OSI layer 1, decides in real-time which users can transmit when and on which channel. Some data link layer protocols (such as those used in Ethernet and cable modems) take a *laissez-faire* approach. They let users attempt to transmit whenever they don't hear any other user talking. Occasionally multiple users will attempt to transmit at the same time; upon learning of the resulting packet collision, these users will have to attempt their transmissions again later. Other data link layer protocols (such as those used in DSL and telephone systems) take a more organized approach. They make users take turns and/or they assign different users to different channels. A few data link layer protocols (such as those used in 802.11 and many cell phone networks) adopt an in-between approach by giving users more direction than the *laissez-faire* ones but less direction than the organized ones.

Anytime that the modulation scheme and the data link layer protocol allow multiple signals to overlap in time, space, and frequency, this phenomenon is called *interference*. In wired networks, interference arises when multiple users or devices *on the same wire* transmit at the same time on the same frequency. In wireless networks, interference is caused when multiple users or devices *within hearing distance of each other* transmit at the same time on the same frequency. Interference in wireless networks is usually worse than in wired networks, because there are typically more wireless devices within hearing distance of each other than wired devices share a wire. As a result, the data link layer protocol is also carefully matched to the type of network, including whether it is wired or wireless.

We now move up to OSI layer 3, the network layer. Portions of the design of OSI layer 3 protocol of a network may depend on whether the network is wired or wireless. In both types of networks, the network layer specifies how devices addresses (e.g. IP addresses) are interpreted, and determines how a packet is routed from one network device to the next. Wireless networks, however, have additional tasks at this layer. They often must consider the allocation of wireless network resources amongst cells (part of *Radio Resource Management*), and this function is implemented in the network layer. In addition, wireless networks with mobile users must accommodate users who move from one cell to another in the middle of a call

¹Higher layers determine how much capacity each user or application is allocated.

or connection (called *handoffs*). In wired networks and in some wireless networks (including Wi-Fi), the Internet Protocol (IP) and its associated routing algorithms are the dominant protocols at the network layer. In some wireless networks (including some cellular networks), a different set of protocols (Signaling System 7) emanating from telephone networks are used to implement a different method of addressing and routing. Sometimes the two sets of protocols are used together, so that higher layer Internet protocols can interface via IP to lower layer Signaling System 7 protocols.

Finally we consider OSI layers 4 through 7. These layers are responsible for managing calls or connections, responding to congestion in the core network, authenticating users, presenting data, and interfacing to applications. Although wireless networks have some special requirements at these layers, most of this functionality is common between wired and wireless networks. There are two dominant sets of protocols used. Most wired and some wireless networks use the Internet protocols, including the Transmission Control Protocol (TCP). Some wireless networks use the telephone network protocols (Signaling System 7). Sometimes variants of Internet protocols are created for wireless implementation, e.g. variants of TCP to cope with the nature of wireless attenuation and multipath.

In conclusion, wireless networks differ substantially from wired networks at the network layer and below. However, they differ in much more limited manners at transport layer and above. In the next section, we turn more specifically to how wired and wireless networks differ with respect to traffic management.

3. HOW IS TRAFFIC MANAGEMENT IN WIRELESS NETWORKS DIFFERENT THAN IN WIRED NETWORKS?

As discussed in section 1, we assume that the primary concern of net neutrality is whether there exists a level playing field between application providers who are also ISPs and those who are not, and that the primary method of potential discrimination is the use of QoS mechanisms such as packet prioritization or bandwidth reservation. The question thus becomes: do the traffic management requirements of wireless networks vis a vis wired networks merit different treatment with respect to net neutrality?

The traffic management requirements of wireless networks can differ from those of wired networks for two broad reasons. First, as discussed in the previous section, wireless networks face greater challenges due to the nature of the wireless medium and often due to mobility of their users. These challenges include attenuation, multipath, interference, and handoffs.

Second, many wireless networks (especially cell phone networks) rely heavily upon their ability to offer satisfactory performance for telephone calls. This reliance is historical. Cell phone networks (as with wired telephone networks) were initially designed to support telephone calls. In contrast, cable tv networks were initially designed to support one-way distribution of video, and the Internet was initially designed to support email and file transfer. Each network's initial architecture reflected the dominant application that consumers of that network paid for.

Each of these three sets of applications also defines good performance in a different manner. Real-time applications such as telephone calls and video conferencing

require a very short end-to-end delay. Performance is measured by the percentage of packets that don't arrive by a fixed deadline of a few tenths of a second. Audio and video distribution, such as the cable tv and streaming, requires a capacity per channel or stream that does not vary much within a period of a few seconds. Performance is measured by the percentage of packets that don't arrive by a fixed deadline of a few seconds. The delay-insensitive applications that currently dominate the Internet are very tolerant of large variations in delay and capacity. Performance is measured by the average capacity (or *throughput*) they obtain over a period of seconds to minutes.

As a consequence, network architects have designed different traffic management techniques for each of these three sets of applications. Networks that specialize in supporting real-time applications, such as telephone networks and cell phone networks, allow only very limited queuing in network devices, so that each device adds only very small queuing delays. In addition, these networks reserve capacity for each call or connection or otherwise limit the total traffic in the network, so that the negative impact of users upon each other is controlled and limited. Networks that specialize in supporting video distribution, such as cable tv, allocate fixed capacities to each channel, so that variability is limited. Networks that focus on supporting delay-insensitive applications, such as the Internet, allow much more queuing in network devices and accommodate much greater variability in throughput per connection, so that a wide variety of applications can be efficiently supported.

In wireless networks, capacity reservation for voice calls is accomplished using a few different techniques. First, the power used to transmit the signal from a consumer's device and vice-versa is adjusted frequently (often many times per second) on the basis of attenuation and multipath. The goal of these adjustments is to maintain a constant quality in the connection. Second, the network limits the number of number of voice users in each cell in order to maintain a minimum performance level per user. When a user migrates from one cell to another, the network attempts to allocate resources in the next cell; if resources are not available in the next cell and if the user requires too much transmission power to maintain a connection with the current cell (thus causing too much interference for its neighbors), the call is terminated. None of these techniques are yet common in the wired Internet.

These wireless traffic management techniques were designed for the dominant application on cell phone networks - telephone calls. However, there is a long term trend toward technology convergence. Both texting and Internet access are now key applications supported by cell phone networks, and subscribers pay significant sums for each. Cell phone networks are thus migrating in their architecture to look more like the Internet. In parallel, on the Internet real-time applications such as telephone calls, video conferencing, and gaming are growing in popularity. In response, Internet architecture is expected to migrate in its architecture to more efficiently support these applications. Cable tv networks now support both Internet access and telephone calls. Its architecture is also migrating toward one that can gracefully support all three sets of applications.

In order to support texting and Internet access on cell phone networks, and

to create wireless local area network protocols such as Wi-Fi, traffic management techniques were required that are appropriate for delay-insensitive applications on wireless networks. These techniques have some commonality with traffic management techniques used in wired networks for delay-insensitive applications, but they must also cope with the variability of the wireless medium. Thus, they also borrow ideas from traffic management techniques used in wireless networks for real-time applications. Like their wired network brethren, wireless traffic management for delay-insensitive applications allows for a wide variety of throughput per connection, so that a wide variety of applications can be efficiently supported. This is often accomplished using scheduling in the data link layer protocol (e.g. determining who can transmit when). Like their wireless real-time application brethren, wireless traffic management for delay-insensitive applications dynamically adjusts transmission power on the basis of attenuation and multipath, and may use some elements of Radio Resource Management to balance load between cells and to support handoffs.

It is thus true that wireless networks require some different types of traffic management than wired networks – due to both the nature of the wireless medium and the greater focus of some wireless networks on real-time applications. In wireless networks, users compete for resources with other users within hearing range, rather than only on the same wire. The greater variability of signals transmitted wirelessly instead of on wires would render real-time applications useless without some type of QoS implemented to smooth out the variations. A portion of these QoS techniques must be applied in very quick response to variations in the wireless signal, and thus must be automated in the wireless device and in network equipment.

In addition, wireless network capacity is usually more expensive than wired network capacity. The cost of wired networks is usually dominated by the cost of purchasing and installing transmission lines, and to a lesser extent by the cost of network devices such as routers. Wireless networks incur similar costs, with the cost of transmission lines replaced by the cost of obtaining spectrum. In addition, wireless networks may incur costs for leasing sites for base stations. However, the capacity of wireless networks is almost always significantly less than the capacity of wired networks that cover the same geography, due to the relative scarcity of spectrum. This decreased capacity translates into a higher shadow cost for bandwidth². As a result, there is often an incentive for traffic management techniques in wireless networks to be more efficient than in wired networks.

In the next section, we turn more specifically to how these differences impact the issue of net neutrality.

4. HOW DO THE DIFFERENCES IN TRAFFIC MANAGEMENT AFFECT NET NEUTRALITY?

Proponents of net neutrality (generally, application providers and consumer groups) argue that without a prohibition on discrimination, ISPs may charge application providers discriminatory prices for access to dedicated bandwidth or for QoS, or may

²Note, however, that in low density areas wireless networks may be lower cost than wired networks, due to the ability of a single wireless base station to replace a large number of long transmission lines.

outright block access to certain applications or websites, and that such activity will inhibit development of new Internet applications (see e.g. Center for Democracy & Technology [2010], Free Press [2010], and Google Inc. [2010]). Some proponents believe that ISPs should not be allowed to charge for priority treatment of traffic on the Internet portion of their service offerings (see e.g. Free Press [2010]). When applying their position to wireless networks, many proponents of net neutrality argue that wireless networks face the same dangers as wired networks (see e.g. Free Press [2010] and Google Inc. [2010]). They are particularly concerned when wireless ISPs restrict the applications used on wireless devices (especially voice and video over IP). Many proponents believe that wireless providers should not restrict applications and should treat all applications equally (see e.g. Free Press [2010], Google Inc. [2010], and Center for Democracy & Technology [2010]).

Opponents of net neutrality (generally, ISPs) argue that there is no current problem, that competition is sufficient to ensure that commercially negotiated arrangements for bandwidth or QoS will not negatively impact consumers, and that any regulation will discourage investment in network infrastructure (see e.g. AT&T Inc. [2010] and Verizon and Verizon Wireless [2010]). When applying their position to wireless networks, opponents of net neutrality see few differences. Opponents usually believe that there is a greater need for traffic management on wireless networks than on wired networks, and that this further undermines the case for wireless net neutrality (see e.g. AT&T Inc. [2010] and CTIA - The Wireless Association [2010]). In addition, they often argue that there is greater competition amongst wireless providers, and thus even less need for net neutrality regulation (see e.g. CTIA - The Wireless Association [2010] and Verizon and Verizon Wireless [2010]).

Some of the questions related to net neutrality take on additional importance in wireless networks. Do users have the right to run any software of their choice on wireless devices? Do they have the right to attach wireless devices of their choice to the network?

Most proponents of wireless net neutrality believe that users should have the right to attach any device and to run any application, so long as they do not cause harmful interference to other wireless users (see e.g. Free Press [2010], Google Inc. [2010], and New America Foundation et al. [2010]). Most opponents of wireless net neutrality believe that an ISP can reasonably dictate which wireless devices can be used on their networks, and that either an ISP or a device manufacturer can reasonably limit which applications can be run on wireless devices; they believe that there is sufficient competition to ensure that social welfare is maximized (see e.g. AT&T Inc. [2010] and CTIA - The Wireless Association [2010]).

In the previous section, we found that wireless networks require some different types of traffic management than wired networks. These differences included: quick response to wireless signal variation, limits on the interference users cause each other, limits on the number of active real-time users, scheduling of transmissions, and reservation or priority of resources for real-time applications. In addition, wireless network capacity is usually more expensive than wired network capacity. The question is: do these differences in traffic management justify differences in net neutrality?

For proponents, these differences pose a conundrum. They have difficulty defin-

ing *harmful interference*, and in pinpointing which traffic management practices are acceptable ways for a wireless network provider to limit harmful interference. More fundamentally, they understand that wireless telephone calls require these stronger types of traffic management, and that this includes some type of bandwidth reservation or priority. However, they wish to apply net neutrality to the remainder of the applications used. As a consequence, most proponents are left with no choice other than to somehow segment off wireless telephone calls from other applications, and to apply net neutrality only to the other applications. One such segmentation is to apply wireless net neutrality only to the *Internet* portion of a provider's offerings. Another segmentation is to apply wireless net neutrality to everything, but then to make an exception for *managed services*. Both approaches face substantial difficulties as technology convergence erases the difference between telephone networks and the Internet. How do you define the *Internet* portion when all applications, including voice and video, run over IP? Where do you draw the line between *managed services* and unmanaged services when there is a wide variety of applications that require a wide variety of QoS?

For opponents, the differences between wired and wireless traffic management are less worrisome, since the stronger techniques used in wireless seem to support their case more strongly. However, opponents have difficulty explaining why wireless traffic management requires limits on devices and applications, when such limits on wired networks would be considered unacceptable. If there is greater competition amongst wireless providers, why doesn't that translate into greater freedom, not less? Opponents often argue that the limits are justified by wireless interference. However, in both wired and wireless networks users compete with each other for resources; why does wireless interference justify limits on applications that wired interference doesn't? Alternatively, opponents argue that the limits are justified by the higher cost of wireless resources. However basic economics would lead one to expect that higher resource costs should translate into higher (or differentiated) service costs, not to somewhat arbitrary limits on use.

Thus, for both proponents and opponents of net neutrality, the differences between traffic management in wired and wireless networks pose a challenge to their positions. Having concluded in the previous section that wireless networks require some different types of traffic management than wired networks, we turn to the question of how the differences in traffic management affect net neutrality. Clearly, despite the differences in traffic management between wired and wireless networks, similar net neutrality concerns apply.

A key finding in our analysis above is that wireless networks differ substantially from wired networks at the network layer and below, and subsequently the differences between wired and wireless traffic management techniques lie almost entirely at or below the network layer, namely at OSI layers 1-3. Protocols that respond to wireless signal variation, that limit interference, that limit active real-time users, that schedule transmissions, and that reserve or prioritize resources, are all in these layers. This is consistent with the Internet layered model, in which these layers are assumed to be carefully matched to the particular transmission medium.

As a consequence, we conclude that wireless networks often are justified in using different traffic management techniques, but only at these lower layers. It follows

that these differences merit a definition of *reasonable network management* that recognizes the differences in lower layer traffic management requirements and techniques. However, it also follows that lack of differences in upper layers merits a definition of *reasonable network management* that enables similar competition at the application layer independent of the type of network. Wireless networks should address their greater challenges at lower layers either by exercising stronger lower layer traffic management techniques than used by wired networks or by exercising reasonable techniques used in wired networks but to a greater extent. However, wireless networks are not justified by technical differences in implementing different traffic management practices above the network layer than those used in wired networks.

We conclude that any net neutrality position applied to wireless networks, whether pro or con, should reflect the differences between wired and wireless networks at or below the network layer, and should reflect the similarities between wired and wireless networks above the network layer. Proponents of wireless net neutrality should accept that wireless networks require stronger forms of traffic management at or below the network layer, and should focus on regulation that ensures a level playing field between providers of various applications (including those that require QoS). Opponents of wireless net neutrality should accept that stronger traffic management at or below the network layer does not justify different treatment above the network layer, and should focus on defining reasonable network management in a manner that acknowledges this. Both sides should accept that segmentation of network applications between managed and unmanaged services, or between Internet and non-Internet offerings, is problematic as technical convergence nullifies such distinctions.

We conclude that wireless networks often are justified in using different traffic management techniques, but only at these lower layers. However, since such lower layer techniques include the reservation and prioritization methods that generated much of the initial net neutrality debate, it remains a challenge to construct a net neutrality policy that can be consistently applied to both wired and wireless networks. We turn to this challenge in the remaining sections of the paper.

5. WIRELESS NET NEUTRALITY VERSUS WIRED NET NEUTRALITY

Up to this point in the paper, we have addressed the differences between wired and wireless net neutrality without taking a position on the issue of net neutrality itself. In the the remaining sections of the paper, we consider potential methods to extend wired net neutrality to wireless networks. In these sections, we do take a position. In this section, we first give our general position on net neutrality, and then consider how it may or may not apply to wireless net neutrality.

In Jordan [2007], we argued for middle-ground on net neutrality. We proposed a policy that (i) bans discrimination on the basis of source, destination, or ownership of traffic, (ii) allows ISPs to implement other QoS mechanisms, (iii) requires that any such QoS mechanisms implemented in network infrastructure be made available without unreasonable discrimination to competing application providers and peering ISPs, (iv) allows ISPs to charge subscribers, application providers, and/or peering ISPs for use of such QoS mechanisms provided these charges are not unrea-

sonably discriminatory, (v) bans unfair methods of competition and unreasonable discrimination, and (vi) provides forbearance when sufficient competition exists in network infrastructure.

The policy is based on a framework that distinguishes between discrimination in high barrier-to-entry network infrastructure and in low barrier-to-entry applications. The policy prohibits use of Internet infrastructure to produce an uneven playing field in Internet applications. In this manner, the policy restricts an Internet service provider's ability to discriminate in a manner that extracts oligopoly rents, while simultaneously ensuring that ISPs can use desirable forms of network management. The paper presented statute language to implement the proposed policy.

On a related note, in Jordan [2009], we extended our analysis to consider other types of traffic management that might fall under the umbrella of net neutrality. We proposed that traffic management practices are reasonable if they are implemented at endpoints, are chosen by the user, are based on reasonable application provider payment, or involve providing QoS to traffic chosen by the user. We proposed that traffic management practices implemented in transit nodes without user choice are unreasonable if they block applications or violate the net neutrality policy discussed above (e.g. provide QoS based on source or on unreasonably discriminatory application provider payment). We suggest that QoS based on application can be more effectively implemented by allowing the user to determine the priority of his/her applications, and we suggest that any charges for QoS can be most effectively implemented by integrating them into subscriber contracts and into the Service Level Agreements between ISPs, rather than by charging application providers that are not subscribers.

We acknowledge portions of this middle-ground position may be unacceptable to some net neutrality proponents and some opponents. Some proponents of net neutrality object to allowing ISPs to implement QoS or to allowing them to charge for it. Some opponents of net neutrality object to prohibitions on deep packet inspection or on unreasonably discriminatory charges.

We now turn to the question of whether such policies on wired net neutrality can be extended to wireless net neutrality. In both wired and wireless networks, unreasonably discriminatory use of reservation of bandwidth and/or prioritization of packets can threaten to tilt the playing field between ISPs and application providers or amongst application providers. The immediate danger in wired networks is that cable ISPs who prioritize their own VoIP traffic may deny access to this same prioritization to competing VoIP providers or may ask unreasonably prices for this QoS, and that ILECs who prioritize their own video over IP traffic may deny access to this same prioritization to competing video providers or may ask unreasonably prices for this QoS. We see the same danger in wireless networks, perhaps even to a greater extent due to the more widespread use of QoS.

The question is how to limit these dangers. One option is to ban the use of reservation and prioritization practices. This option is unreasonable, since these QoS practices are often necessary to support real-time applications. A similar option is to restrict their use to *managed services*. This option is impractical, since as discussed above we don't believe one can define such a class; applications change

very rapidly, and it is this flexibility and evolution that must be maintained. Rather than banning or limiting QoS, some proponents of net neutrality would mandate open access to all lower layer components and protocols to application providers. We don't believe this is necessary to accomplish the goals.

Requiring an *open interface* to QoS mechanisms in lower layers is simpler and more efficient than requiring open access or defining *managed services*. The idea comes from separating networking functions into two subsets: OSI layers 1-3 (network infrastructure) and OSI layers 4-7 (network applications). Network infrastructure layers exhibit a high barrier-to-entry and hence have small number of providers, while the application layers exhibit a low barrier-to-entry and hence have a large number of providers. Net neutrality can be effectively achieved by a properly regulated open interface *from applications to network infrastructure*. The requirement of an open interface captures the central tenet of a layered Internet architecture, and yet is less intrusive than a requirement of open access. In particular, it is less intrusive than requiring a standard air interface for wireless devices, which may be required to ensure the right to attach devices; an open interface between the networking and transport layer only requires standardization of the *Application Programming Interface* (API) for services offered by the networking and lower layers, while a standardized air interface requires standardization of all physical and data link layer protocols and of their use.

In Jordan [2007], we proposed that any QoS mechanisms that an ISP implements in network infrastructure layers should be available to application providers without unreasonable discrimination. Requiring such an open interface can ensure that ISPs are prohibited from refusing to provide enabling Internet infrastructure services to competing application providers in order to differentiate the ISP's own application offerings, prohibited from providing Internet infrastructure services to competing application providers at inflated prices in order to favor the ISP's own application offerings, and prohibited from making exclusive deals to provide enabling Internet infrastructure services to certain application providers. It can also ensure that ISPs have the right to apply network management mechanisms that do not threaten a level playing field, and to make arrangements with consumers, application providers, and peering ISPs for Internet infrastructure services in a manner that does not conflict with the above goals. Finally, forbearance should be applied to these regulations where sufficient competition exists.

But is an open interface an effective way to require wireless net neutrality? The differences between wired and wireless traffic management lie almost entirely within the network infrastructure layers. Hence, an open interface would similarly require that the QoS mechanisms used in these wireless layers be available without unreasonable discrimination to application providers. There is no need to define *harmful interference*, as the ways in which users compete for resources with each other continues to be controlled by the network infrastructure layers; an open interface does not impede an ISP's ability to reasonably accomplish this. There is no need to define what constitutes the *Internet* portion of a provider's offerings, as the layering model applies both to Internet protocols and to telephone network protocols; all of a wireless providers offerings are thus included without reference to the application or technology. There is also no need to define what constitutes *managed services*,

as the open interface requires access to lower layer QoS mechanisms that enables real-time applications, rather than carving out real-time applications as an exception; the open interface thus encourages competition in managed services, rather than inhibiting such competition. Finally, requiring an open interface in both wired and wireless networks avoids the need to differentiate net neutrality policy on the basis of the technology used. However, it does allow for differences between wired and wireless networks on the basis of competitiveness, since forbearance may be granted in some markets but not in others.

An open interface also would bear on the question of whether users have the right to run the software of their choice on wireless devices. Because an open interface allows any application provider to offer any functionality within the network application layers without unreasonably impacting the network infrastructure layers, there is no reasonable justification on the basis of traffic congestion for limiting applications on any device. The impact of an application would be controlled by the ISP at the interface through limits and/or charges placed on traffic. Applications that transmit high volumes of traffic would either purchase this capacity at standard rates or would consume a high proportion of a pre-purchased traffic quota. Applications that require QoS would similarly either purchase the required QoS at standard rates or would consume some portion of pre-purchased QoS. The interface is only concerned with traffic volume, QoS, and payment; it is not concerned with what particular application this traffic is destined to. From the point of view of the lower layers, therefore, the system is application-agnostic.

This interface-based approach is consistent with current wireless device and operating system architecture. There are a number of different operating systems that are used on wireless devices. However, they are all built using layered architectures. The operating system limits access of applications to lower layer protocols. Indeed, they provide an interface between applications and the operating system called an *Application Programming Interface* (API). The API defines what services the operating system offers to applications and how to access them. As in wired networks, common wireless network device operating system APIs offer access to lower layer functionalities at the the networking layer or above; they do not offer direct access to protocols at the physical and data link layers. This architecture implies that *open access* to all lower layer protocols is not necessary for net neutrality; all that is required is an *open interface*. Similarly, the architecture implies that lower layer protocols have no need to know what application a packet belongs to; they only need know what type of traffic management to apply.

We turn in the next section to consideration of what such an open interface would look like.

6. HOW CAN PRICING BE USED TO DEFINE AN OPEN INTERFACE?

The goal of an open interface should be to allow application access on a level playing field to lower layer QoS. The interface also represents a contract between the wireless carrier and the user. The wireless carrier offers certain services (perhaps at a specified price) and the user requests (or purchases) those services desired at the volume desired. The central idea is that the wireless carrier will use traffic management techniques to ensure that services are rendered and hence contracts

are satisfied.

Since the services are offered at an interface between the network and transport layers, they can not be based on the application. They can, and likely will, however offer a range of QoS options. Some of these options may be targeted to support particular applications efficiently, e.g. high throughput service for file sharing, traditional best effort service for email and web browsing, bounded delay service for streaming, and guaranteed low delay service for VoIP and video conferencing. However, the choice of which application to transmit over which QoS option is up to the user; if a user wants traditional best effort service for file sharing, she can elect this option.

The terms of the contract at the open interface likely involve some form of pricing. Two forms are currently common. In the quota form, a user can purchase access to a specified amount of a service at a specified price, e.g. 5GB/month of best effort service with a peak transmission rate of 1Mbps for \$60/month. In the volume form, a user can purchase access on a per volume basis, e.g. \$.01/min/(8kbps) for guaranteed low delay service (where 8kbps is chosen to accommodate one VoIP stream). The two form are often combined, e.g. 5GB/month for \$60/month with excess volume charged at \$50/GB³.

The most likely deployments that would be consistent with this approach may be user tiering and/or application provider payment. Under user tiering, a user may purchase a higher tier that includes the ability to transmit and/or receive a certain amount of QoS prioritization, e.g. enough to support up to 1000 minutes of VoIP prioritization (if that is how the subscriber chooses to use it). Under application provider payment, a VoIP provider other than the user's ISP may wish to purchase a VoIP prioritization option on behalf of the user (perhaps bundled with the VoIP service offered to the user); if the payment was reasonable this would be consistent with our approach.

In contrast, many current plans are not application-agnostic and are hence not consistent with an open interface. Some plans for smartphones include unlimited amounts of data, but restrict use to certain devices (e.g. prohibit tethering to a laptop) and to certain applications (e.g. permit web browsing and email, but prohibit file sharing, streaming, and VoIP). The goals of traffic management can be more efficiently obtained through an application-agnostic interface that allows users to choose their own applications and to match these applications to QoS options based on price.

The form of the contract at the open interface and the prices charged will likely relate to usage of resources in lower layers. In wireless networks the key transmission resources are bandwidth and transmission power. The bandwidth and transmission power allocated to a user or a flow largely determines the low level performance experienced. This low level performance is usually described as a combination of throughput, delay, and packet loss. Various lower layer traffic management practices control the bandwidth and transmission power allocated to each user or flow in order to create the desired variety of QoS service offerings. The price for each QoS service offering thus depends in part on the resources used to achieve that QoS.

³These are the charges for one of the AT&T Wireless Data Connect plans at the time of writing.

Finally, it should be noted that placing the contract at the interface between the network and transport layer in no way unreasonably restricts an ISP's ability to charge additional amounts for services that are not part of OSI layers 1-3. An ISP may continue to subsidize the purchase price of a wireless device and to recoup that subsidy over time. An ISP may also offer high layer services at additional cost, e.g. voice mail, voice dialing, navigation, ring tones, and roadside assistance.

7. WIRELESS NET NEUTRALITY

In this final section, we consider the application of potential net neutrality statute language to wireless networks. In Jordan [2007], we proposed statute language designed to require net neutrality through an open interface. The language starts with the following definitions in order to specify the location of the interface:

DEFINITIONS.

(1). *INTERNET INFRASTRUCTURE SERVICES.*-The term 'Internet infrastructure services' means all services- (A) over a network that uses a public right-of-way; and (B) that reside at or below the network layer or are required to manage the network.

(2). *INTERNET APPLICATION SERVICES.*-The term 'Internet application services' means all services-(A) over a network that uses a public right-of-way; (B) that are not infrastructure services; and (C) that do not fall under title VI of the Communications Act.

(3). *NETWORK LAYER.*-The term 'network layer' means the third layer of the 7-layer Open Systems Interconnection Model, responsible for message addressing and for routing information within the network, including routing within the telephone network and including the Internet Protocol within the Internet.

(4). *ACCESS NETWORK* -The term 'access network' means the portions of the Internet service provider's network which must be transversed to form routes from the Internet to its subscribers.

Internet infrastructure services thus define the lower layer services, whereas *Internet application services* define upper layer services. Internet infrastructure services can only be provided by carriers, and must be provided by each carrier on their portion of the network. Internet infrastructure services require large investments and thus exhibit a high barrier-to-entry. Internet application services can be provided by carriers or by many other application providers on the Internet. Such services usually exhibit a low barrier-to-entry. Among Internet infrastructure services, it is those provided on *access networks* that often have limited competition.

As discussed above, the same interface – between the network and transport layers – provides the meaningful distinction in wireless networks as in wired networks. It should also be noted that almost all wireless networks are access networks.

The first key portion of the proposed language prohibits ISPs from refusing to provide enabling Internet infrastructure services to competing application providers in order to differentiate the ISP's own application offerings, from providing Internet infrastructure services to competing application providers at inflated prices in order to favor the ISP's own application offerings, and from making exclusive deals to

provide enabling Internet infrastructure services to certain application providers. This is accomplished by extending the prohibition in title II of the Communications Act on unjust or unreasonable discrimination in charges to cover Internet infrastructure services:

DISCRIMINATION AND PREFERENCES.

For purposes of sections 202 and 206 through 209 of the Act, an Internet service provider shall be treated as a common carrier, and Internet infrastructure service shall be treated as a communications service.

These limits are applied only to services that can be provided only by the ISP, and this does not apply other title II requirements to ISPs, e.g. tariffs, interconnection, or unbundling. Based on the above discussion, these same limits on use of discrimination and preferences should apply to wireless networks as to wired networks.

The next key portion of the proposed language concerns the responsibilities of vertically integrated ISPs to *their own subscribers and peers*. When an ISP offers applications that rely on Internet infrastructure services, ISPs are required to make available to competitors the same Internet infrastructure services at the same prices:

COMPETITION.

An Internet service provider shall make available to subscribers and other Internet service providers - on the same prices, terms, conditions of sale, and delivery - any Internet infrastructure services provided on its access networks as the Internet service provider offers to Internet application services provided by itself or its affiliates. An Internet service provider shall provide Internet infrastructure service to subscribers and other Internet service providers, that is at least equal in quality to that provided by the Internet service provider to itself or its affiliates.

Next we turn to the responsibilities of vertically integrated ISPs to *subscribers of other ISPs*. A vertically integrated ISP is prohibited from using service level agreements with its peers to favor its own applications.

It is unlawful for an Internet service provider to engage in unfair methods of competition, unreasonably discriminatory conduct, or unfair or deceptive acts or practices, the purpose or effect of which is to hinder significantly or to prevent any Internet application provider from providing content, applications, or services to consumers.

Each of these three provisions are general enough so that we see no aspect of wireless networks that would justify different treatment from wired networks.

The next key portion of the proposed language ensures that ISPs have the right to apply network management mechanisms that do not threaten a level playing field, and that regulation does not impede an ISP from making arrangements with consumers, application providers, and peering ISPs for Internet infrastructure services in a manner that does not conflict with the other goals. To accomplish this, in addition to a number of protected management techniques that had been previously proposed, we suggested adding provisions to explicitly guarantee ISPs the right to alleviate congestion by treating all traffic similarly or by treating all applications of

the same type similarly, to sell reserved bandwidth and QoS to both their residential and business subscribers, and to discriminate in the carriage of Internet traffic based on peering arrangements with other ISPs:

NETWORK MANAGEMENT.

An Internet service provider may alleviate congestion in a manner that does not distinguish based on the source or ownership of content, application, or service, may offer directly to a subscriber Internet service at different prices based on defined levels of bandwidth, quality of service, or the actual quantity of data flow over a user's connection, and discriminate in the carriage and treatment of Internet traffic based on such contract with that subscriber; and may enter into contracts with other Internet service providers, and discriminate in the carriage and treatment of Internet traffic based on such contract with that Internet service provider.

These provisions should also apply to wireless networks. Wireless network ISPs may use different methods to control congestion, and may exercise these methods more strongly due to more limited resources. They may also offer QoS directly to subscribers for a price and/or may arrange with other ISPs to honor QoS agreements made by them. None of these practices would threaten a level playing field so long as they are done without unreasonable discrimination.

The final key portion of the proposed language limits application of all of the previous provisions where sufficient competition exists. This is accomplished by ensuring that the forbearance provisions currently in title I of the Communications Act apply to ISPs offering Internet infrastructure service:

COMPETITION.

For purposes of section 10 of the Act, an Internet service provider shall be treated as a telecommunications carrier, and Internet infrastructure service shall be treated as a telecommunications service.

This provision should also apply to wireless networks. Indeed, wireless networks have the potential to exhibit sufficient competition so that net neutrality may not be required. In addition, if these wireless networks are treated comparably to wired networks and offer similar services, then this competition may allow forbearance from net neutrality provisions for competing wired networks.

8. CONCLUSION

The principal question addressed in this paper is whether differences between wired and wireless network technology merit different treatment with respect to net neutrality. Unlike other papers in the academic literature, we focus on applications and traffic management, rather than device attachment. We found that wireless networks differ substantially from wired networks at the network layer and below, and that wireless networks often require different traffic management practices at these lower layers. However, since the differences are confined to these lower layers, we argued that net neutrality in both wired and wireless networks can be effectively accomplished by requiring an open interface between network and transport layers.

Furthermore, we believe that this approach is a more streamlined and more effective solution that carving out a set of managed services. We did not address here several important aspects – in particular, we did not consider potential differences in competitiveness (which may lead to forbearance of our proposed net neutrality rules), and we did not consider the related issue of device attachment.

REFERENCES

- AT&T INC. 2010. Comments of AT&T Inc. before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- BRADEN, R. 1989. Requirements for Internet hosts – communication layers. Tech. rep., IETF RFC 1122.
- CENTER FOR DEMOCRACY & TECHNOLOGY. 2010. Comments of the Center for Democracy & Technology before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- CLARK, D. D. 2007. Network neutrality: Words of power and 800-pound gorillas. *International Journal of Communication* 1, 701–708.
- CTIA - THE WIRELESS ASSOCIATION. 2010. Comments of CTIA - The Wireless Association before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- FCC. 1968. FCC 68-661, Carterfone Order.
- FCC. 2005. FCC 05-151, Internet Policy Statement. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.pdf.
- FCC. 2009. FCC 09-93, Open Internet NPRM. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-09-93A1.pdf.
- FREE PRESS. 2010. Comments of Free Press before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- GOOGLE INC. 2010. Comments of Google Inc. before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- HAHN, R. W. AND LITAN, R. E. 2007. The myth of network neutrality and what we should do about it. *International Journal of Communication* 1, 595–606.
- HAHN, R. W., LITAN, R. E., AND SINGER, H. J. 2007. The economics of "wireless net neutrality". *Journal of competition law and economics* 3, 399–451.
- JORDAN, S. 2007. A layered network approach to net neutrality. *International Journal of Communication* 1, 427–460.
- JORDAN, S. 2009. Some traffic management practices are unreasonable. In *Proceedings of the IEEE International Conference on Computer Communications and Networks (ICCCN)*. 137–140.
- NEW AMERICA FOUNDATION, COLUMBIA TELECOMMUNICATIONS CORPORATION, CONSUMERS UNION, MEDIA ACCESS PROJECT, AND PUBLIC KNOWLEDGE. 2010. Comments of New America Foundation, Columbia Telecommunications Corporation, Consumers Union, Media Access Project, and Public Knowledge before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- PEHA, J. M. 2007. The benefits and risks of mandating network neutrality, and the quest for a balanced policy. *International Journal of Communication* 1, 644–668.
- SKYPE COMMUNICATIONS. 2007. A petition before the Federal Communications Commission to confirm a consumer's right to use Internet communications software and attach devices to wireless networks.
- VERIZON AND VERIZON WIRELESS. 2010. Comments of Verizon and Verizon Wireless before the Federal Communications Commission in the matter of preserving the open Internet and in the matter of broadband industry practices.
- WU, T. 2007. Wireless Carterfone. *International Journal of Communication* 1, 389–426.